Analysis of Human Grasping Behavior: Correlating Tasks, Objects and Grasps

Thomas Feix, *Member, IEEE*, Ian M. Bullock, *Student Member, IEEE* and Aaron M. Dollar, *Member, IEEE*

Abstract—This paper is the second in a two-part series analyzing human grasping behavior during a wide range of unstructured tasks. It investigates the tasks performed during the daily work of two housekeepers and two machinists and correlates grasp type and object properties with the attributes of the tasks being performed. The task or activity is classified according to the force required, the degrees of freedom, and the functional task type. We found that 46 percent of tasks are constrained, where the manipulated object is not allowed to move in a full six degrees of freedom. Analyzing the interrelationships between the grasp, object, and task data show that the best predictors of the grasp type are object size, task constraints, and object mass. Using these attributes, the grasp type can be predicted with 47 percent accuracy. Those parameters likely make useful heuristics for grasp planning systems. The results further suggest the common sub-categorization of grasps into power, intermediate, and precision categories may not be appropriate, indicating that grasps are generally more multi-functional than previously thought. We find large and heavy objects are grasped with a power grasp, but small and lightweight objects are not necessarily grasped with precision grasps—even with grasped object size less than 2 cm and mass less than 20 g, precision grasps are only used 61 percent of the time. These results have important implications for robotic hand design and grasp planners, since it appears while power grasps are frequently used for heavy objects, they can still be quite practical for small, lightweight objects.

Index Terms—Human grasping, manipulation, activities of daily living, prosthetics, robotic hands

1 Introduction

Due to the complexity of the human hand and the wide variety of movements it is able to accomplish, many of the factors influencing the human grasp choice are still poorly understood. There has been considerable effort in determining how certain parameters influence human grasping. In general, it is assumed that there are at least two important factors influencing the grasp kinematics: the object being manipulated and the task to be accomplished [1], [2]. A better understanding of human grasping behavior can be used to provide performance specifications and test cases for robotic systems which are designed to operate in human environments. It can provide basic useful heuristics for grasp planning systems and can define the most essential hand functions which should be restored during rehabilitation. Finally, it can help in designing devices that interact with the human hand, to ensure that they will fit into natural hand use patterns, and to better predict how they will be grasped.

In the accompanying paper [3] we focus on the properties of objects being grasped during the daily work of two housekeepers and two professional machinists, filmed through a head-mounted wide-angle video camera, and correlating those properties to the choice of grasp used with them. In this paper, we now turn our attention to the properties of the task being performed in those videos. The goal of the first part of this paper is to assign to each task distinct properties that are hypothesized to influence the grasp. We then use our extensive video analysis [4] to investigate the influence of the task properties on grasp choice. The second part of the paper correlates the relationship between task, object properties, and grasp choice, with a focus on investigating of how strongly different properties of the task and object influence grasp choice and how well those properties can predict the grasp chosen. The data set used in this publication and in the related publications [3], [4] was made public and can be downloaded [5].

In terms of related work, a number of studies have investigated various manipulative task performances in daily life. In the rehabilitation literature, many studies discuss the "Activities of Daily Living" (ADLs) [6], which allow a person's level of independence/impairment to be assessed. However, ADLs are not well structured to describing the details of hand usage, and therefore the specific problems related to hand use are usually not identified [7]. In one relevant study that investigated objects and ADLs (but not grasps), objects were marked with radio-frequency-identification (RFID) tags and it was shown that 88 percent of the 14 ADLs studied could be predicted correctly by looking at the sequence of handled objects [8], supporting the assumption that the task and the object are closely related. Another study that gives some insight into the most common human manipulation tasks used the International Classifications for Functioning, Disability and Health (ICF) issued by the World Health Organization (WHO) [9] and calculated the frequency distributions of activities [10]. Overall, 3,964 activities of a single healthy person were recorded, with the most common activities shown to be "lifting", "putting down objects", "preparing complex meals", "fine hand use, other specified". Time

- *The authors are with the Department of Mechanical Engineering and Materials Science, Yale University, New Haven, CT. E-mail: {thomas.feix, ian.bullock, aaron.dollar}@yale.edu.*

Manuscript received 12 Aug. 2013; revised 7 May 2014; accepted 19 May 2014. Date of publication 28 May 2014; date of current version 15 Dec. 2014. Recommended for acceptance by A. Okamura.
For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.
Digital Object Identifier no. 10.1109/TOH.2014.2326867

usage research has also tried to determine what tasks and activities the human engages over the course of the day [11], [12], but the long time increments used resulted in only very general results.

In terms of correlating grasps to tasks, one study by Klatzky et al. looked at the connection between cognitive object representations and hand shape categories [13]. It showed that there exists a relationship between the hand shape and the high level object class, and also highlighted that those hand shapes are modulated by the task, as defined via three functional classes (which we use in the task classification in this paper). A number of additional studies have investigated how task constraints influence grasp kinematics through laboratory-based experiments typically involving a very specific sequence of action. The hand kinematics are recorded and how they are modulated with respect to the specific goal of the study are noted. These experiments have shown that the grasp kinematics can be influenced by a task later on in the action sequence [14], [15] as well as the intention upon which the object is picked up [16], [17], [18], [19]. Even though these studies clearly show how the human grasp kinematics are affected by certain tasks properties, they are limited to the laboratory and the specific movements studied.

Related to the prior work described above, there is extensive evidence that the task or intended activity influences the grasp kinematics. However, many challenges remain in terms of defining specific relationships between the two in order to analyze their connection.

The remainder of this paper is structured as follows. In the next Section the properties of the task classification are presented. Section 3 provides an overview of the data acquisition and how the assignment of the task properties was done for our real life data set. Following that, Section 4 describes the results of the task classification. In Section 5, the data from the task and the object classification are combined and their results are presented. Finally, Section 6 discusses the task and object-task-grasp results and Section 7 concludes the analysis.

2 Task Classification for Grasping

2.1 Scope

Our goal for a task classification is to assign properties to the tasks that we believe directly relate to how the human chooses grasps while still retaining a level of generality that allows different tasks to be compared. Specific task descriptions such as "hammering", "using a screwdriver", or "wiping" might be used, but these task descriptions, while intuitive, cannot be compared easily without breaking down each task further. Therefore we seek to assign to each task parameters that allow comparing various attributes between tasks.

In this work we define the extent of a task by a single grasp-release cycle. Since the chosen task properties are general, they can be assigned to each grasp-release cycle. They do not depend on verbal descriptions that cannot be abstracted and compared easily between. Although each grasp-release cycle may be part of a longer high level task, these higher level activities are not analyzed in this paper.

2.2 Task Properties

We assign each task a set of properties that we hypothesize are important for grasping and that can also be easily assigned by visual observation. After a thorough review of the literature and consideration of the wide range of possibilities, we decided to classify tasks along three dimensions: constraints, which describe the degrees of freedom and nature of the constraints on the grasped object; functional class, which describes at a high-level what is being done with the object; and force, which describes whether the force being applied is to lift the object or based on some additional task property such as opening a door.

The first property is the *constraints* of the task. Depending on the task (and also the object properties), an object is only allowed to translate and rotate in certain directions in order to successfully complete the task. We follow the definitions from [20], [21], which specify a complete set of 20 possible relative motions between two rigid bodies. The nomenclature defines the relationship between the object and the environment (a fixed reference frame). Fig. 1 (adapted from [18]) explains the nomenclature and gives examples of human manipulation tasks falling into that category. For example, consider the task of opening a door. In this case there are five constraints on the movement, with only a single possible rotation around the hinge axis of the door. Thus, the constraint assigned is <u>r</u>otation, fi<u>x</u>ed, fi<u>x</u>ed ("rxx") and the object's physical restrictions determine the constraint. However, consider the constrained task of holding a fully filled glass of water. Although the glass is physically free to move, to successfully complete the task, the glass has to be kept upright. Too much tilting would result in spilled water, thus a failure in task completion. We assign the holding a glass task a <u>u</u>nconstrained, <u>t</u>ranslation, <u>t</u>ranslation ("utt") constraint. In this case the task itself defines the constraint, rather than a physical connection to the environment. In a majority of real world tasks the object moves to some extent in all six DoF, but in some directions the movements are negligible. In the holding a glass example, tilting the glass a miniscule amount will be allowable, however the movement amplitude is much smaller as compared to the directions in which it is allowed to move. These small deviations are disregarded and a higher number of constraints is assigned, since the goal is to capture the essential task constraints.

The *functional class* is a high level description of the task, for which we use the three categories from [13]. The "hold and pick up" (referenced by "hold" in the paper) category defines general manipulation activities, primarily object transport. The second category is "use", where the object is used according to the object's specific purpose it is designed for or is commonly used. For example, wiping with a sponge or writing with a pencil are within the "use" category. Transporting a pencil, however, would be classified as "hold and pick up", since it is outside the specific writing tasks a pencil is designed for. The last category is "feel and touch" (referenced by "feel" in the paper), in which the hand is used as a sensory tool to interact with the environment. In addition, cases where contact is made without a specific goal are assigned this category.

The *force* property specifies what type of force is necessary to complete the task. Since the forces required can be

Constraints	Example Image	Class	Example
0		uuu	Free motion
1	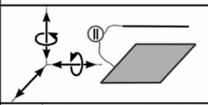	uut	Line parallel to plane
1	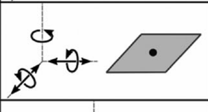	uur	Point on plane: vaccum, writing
2	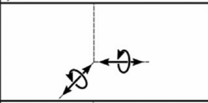	uux	?
2	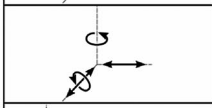	utr	Edge against surface
2	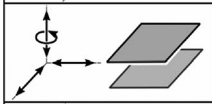	utt	Surface parallel to surface: holding glass, pointing
2	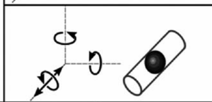	urr	Sphere in slot
3	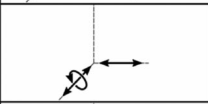	utx	?
3	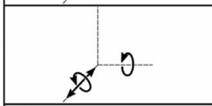	urx	?
3	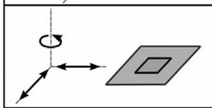	ttr	Surface against surface: wiping
3	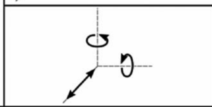	trr	Round peg in slot
3	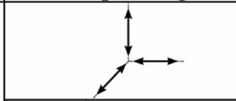	ttt	Only translation
3	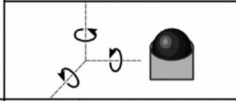	rrr	Ball in socket: ball joint
4	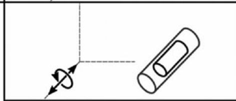	uxx	Cylinder in slot
4	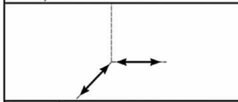	ttx	2D translation
4	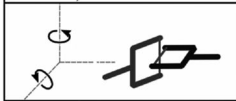	rrx	Cardan/universal joint
4	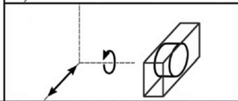	trx	X-slider in slot: sawing
5	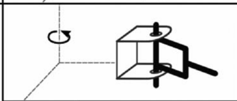	rxx	One rotational DoF: Door
5	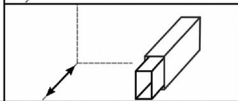	txx	One translational DoF: Drawer
6		xxx	Fixed in space

Legend:

Symbol	Translation/Rotation	Interpretation
u	unconstrained/unconstrained	unconstrained
t	unconstrained/fixed	translation
r	fixed/unconstrained	rotation
x	fixed/fixed	fixed

Fig. 1. Constraint definition for the tasks, adapted from [20]. There are twenty possible classes of object motion relative to the world reference frame. We follow the definition of [20], [21], however change the letters used. Each of the three axes can either be free to move (u), only allow translation (t), only allow rotation (r) or does not allow any movement around that axis (x). Switching two axes within one category results in the same constraint type; letters are ordered in order to facilitate comparison.

complex and difficult to discern visually, we use a simplified description that still provides useful information about the task. Specifically, we assign a value of either "weight" or "interaction". We assign "weight" if the grasp force is closely related to lifting the object. This can be the case for tasks other than object transport, such as using a drill. In that case, the dominating force requirement is to lift the drill, squeezing the trigger usually needs less force. In the second category, "interaction," the grasp force is determined by factors other than object weight, usually through the interaction with the environment. There are two main mechanisms for this decoupling: the weight of the object is supported by the constraints, making the force needed to move the object less than would be required to lift the object (such as opening a drawer or door); or when the interaction force is primarily intended to apply a force through the object, such as is done when scrubbing with a sponge (where the force needed to lift the sponge is much less than the force needed to scrub effectively).

2.3 Limitations of the Classification

In general, defining task properties is much more challenging than assigning object properties, given the wide range of possible classifications. While there are still other parameters of a task which affect grasp choice, many of them are hard to categorize and difficult to assign consistently. Most intuitive task parameters are qualitative properties, which do not generalize well to a physical property. This fact is reflected in a lower inter-rater agreement (see Section 3.1) compared to the object classification. This is also the reason we reduced the task classification to a few essential properties, and removed some properties we tested in the beginning.

We initially assigned a *precision* parameter related to how accurately the movement must be conducted, which we allowed to be either "rough", "normal" or "precise". This parameter might appear very straightforward and easy to understand. However, we found that the raters disagreed (Kohen's $\kappa = 0.09$) for many of the tasks, thus this parameter was excluded from the final data. We believe precision is hard to estimate from visual information in part because it is a general term, but further breaking it up into the precision of individual motion or force components would be too difficult for the current study.

For the force property, the "interaction" category is very diverse, simply stating that the grasp force is not directly determined by object weight. Although additional classification of task forces could help predict grasp choice, in practice it would be difficult to visually classify the task force in much more detail than is already done.

3 Methods

The same data set and approach as in the accompanying paper [3] was used. The data is based on two housekeepers and two machinists, recorded during their professional work for approximately 8 hours each [4], [22]. In the initial video processing only high level task names were recorded, which were classified in a second step. Similar to the object classification, we allow raters to set the task to "cannot classify" to accommodate tasks which cannot be reliably classified based on the high level task description and video snapshots provided to the raters.

3.1 Inter-Rater Agreement and Error Estimation

To evaluate the reliability of our classifications, we took the same approach as in our joint paper, Section. 3.3 [3]. All 297 tasks were classified by two raters based on the category descriptions laid out in Section 2.

They raters agreed in 82 percent (Cohen's $\kappa = 0.16$ [23]) of the instances whether the task can be classified or not. Rater 1 assigned "cannot classify" rarely (10 tasks, 3 percent), whereas Rater 2 assigned it in 58 (20 percent) of the tasks. This explains the large discrepancy between the agreement and the Cohens Kappa. To further analyze the agreement between raters, instances where at least one rater assigned cannot classify were removed. This allows to compare the assignment from both raters. The low agreement on whether the tasks can be classified does not influence the following analysis as data from both raters is present.

For the force parameter, the raters assigned the same property in 84 percent (Cohen's $\kappa = 0.68$) of the instances.

Concerning the constraints, the raters agreed in 68 percent of the cases (Cohen's $\kappa = 0.59$). The main difference is that rater 1 assigned "uuu" (fully unconstrained) in some cases where rater 2 assigned some task constraints. In 12 tasks rater 1 assigned "uuu" and rater 2 assigned "txx" (one translational DoF). Those instances include tasks like "zipping jacket", "removing ruler from pocket" and "moving curtain". Here, it becomes clear that the two raters had slightly different opinions about the nature of the task. The other common discrepancy is "uuu" and "rxx" (one dimensional rotation), which includes tasks like "turning tool" and "turning page". Clearly all those tasks have some one dimensional movement, but also might be complex enough that in reality all six object movements matter. This shows the complexity of a verbal description, it is very challenging to establish a good rule upon which a task is regarded to be constrained.

Finally, agreement on the functional class was found in 66 percent of the instances (Cohen's $\kappa = 0.37$). Rater 2 assigned the "use" property much less; there are 69 cases in which rater 1 assigned "use" and rater 2 assigned "hold". This discrepancy was mainly that rater 2 had a more strict view about what was regarded to be "using and object". As this view was too restricted, the task description was adapted accordingly to make the distinction between "use" and "hold" clearer.

The 95 percent confidence interval of the results was estimated using the same bootstrap method as described in the joint paper, Section 3.4 [3]. This method takes into account both the rater disagreement and the frequency distributions in the data.

3.2 The Final Task Classification Data Set

The classifications of the two raters were combined in a semi-supervised method. The differences in the assignments were analyzed by rater 1, who then made the final classification decision. This step added another layer of review of the assignments, where unclear task descriptions were identified and potential misclassifications were corrected. As described in the inter-rater agreement, there was a discrepancy in the function class that rater 2 assigned "use" more conservative. Most of those differences were corrected to be "use". Overall this step increased the overall accuracy of the data set. The error bar estimation reflects the disagreement of the raters, thus gives a clear estimate on how reliable the results are.

4 Results

4.1 General Statistics

After classifying and rejecting the instances with an associated task that could not be classified, 9,933 grasp instances are used for further analysis. Overall, 231 tasks are present in the data set.

Fig. 2 shows the 12 most frequent tasks and their distribution within the grasp types (less frequent tasks are white, separated by a vertical line), utilizing the task names chosen by the individual raters in the initial video tagging (with a goal of being descriptive within one-three words). There is a good deal of overlap, as can be seen in the multiple instances of "holding", resulting from the fact that the rater tried to best capture what was happening and adding additional context to purely assign "holding" as high level task. For example holding a glass might have different constraints than holding a tool and that would be reflected in the task description. These frequent tasks will have a large influence on the final results. Compared to the object data set, the 12 most frequent tasks cover 72 percent of the data set, whereas the 12 most frequent objects cover only 56 percent [3] of the data set. However, their distribution over grasp types is relatively even - from the most common grasp types only the precision disk is weighted heavily by one task ("wiping surface").

4.2 Overall Task Properties

In this section, we present the results after the taxonomies in Section 2 were applied. Concerning task force, Fig. 3 shows that in 60 percent of the instances the force is related to lifting to object ("weight" category). However, in the

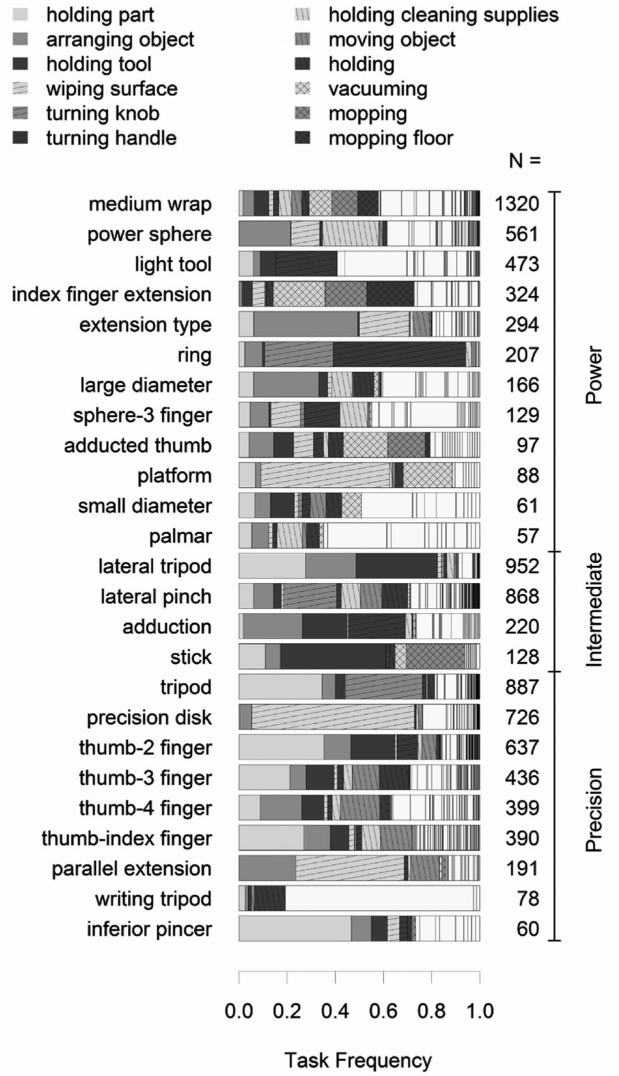

Fig. 2. Task-grasp relationship for common tasks. The most common 12 tasks are highlighted, while all other tasks are white. Some grasps are weighted by only a few very common tasks. Grasps with fewer than 50 instances are omitted. The order of the legend is equal to the order within each bar.

remaining 40 percent of the cases, the mass of the object does not give any information on the grasp force.

No constraints are present (“uuu”) in 54 percent of the instances. There are 10 other constraint configurations present, of which four are more common. The “uur” constraint mainly involves writing, vacuuming and mopping. Another common constraint is “utt”, which is used when pointing in one direction, such as when spraying or holding a glass. For wiping a surface and cases where an object is moved on a plane, the “ttr” property is assigned. Finally, “rxx” involves only a single rotational degree of freedom, for tasks such as moving a door. This final “rxx” constraint is the most common constraint, found in about 17 percent of the instances.

Functional class is dominated by the “hold” category, which was assigned in 54 percent of the grasps. In about 46 percent of the instances the object was “used” and in only 9 total instances (less than 0.1 percent) the functional class was “feel”.

Fig. 4 shows the correlations between the three task categories within the data set. The most prominent parameter

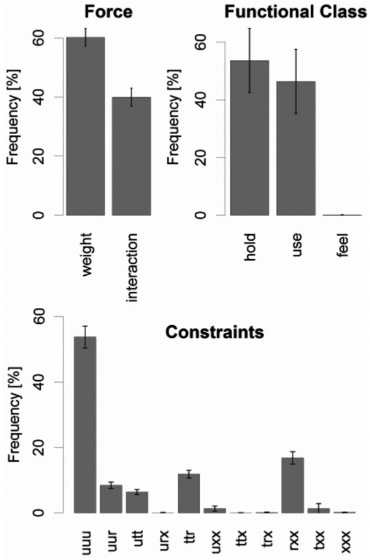

Fig. 3. Distribution of the task properties. The error bars represent a 95 percent confidence interval, as described in Section 3.1.

combination is “hold”, “weight” and “uuu”. This combination accounts for 5,042 instances (51 percent of the data), and generally applies to unconstrained object transport tasks. The other common combinations are within “force = interaction” and “use”, which involves a much wider variety of constraints.

Considering the force-constraint relationship (Fig. 4), shows that when constraints are present, the grasp force is usually decoupled from the weight of the object. The only big exception is the “utt” case, which is mainly used for pointing a spray bottle onto a surface, a frequent

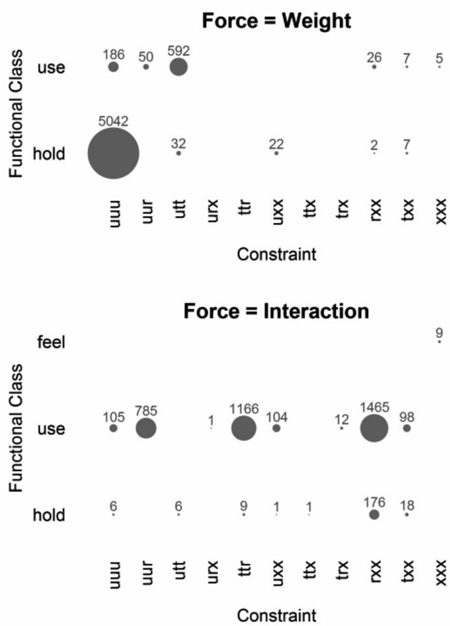

Fig. 4. The two plots show the correlations within the data set. The top plot shows the correlations for all datapoints where the force equal to weight and the second plot shows all instances where the force is not related to the weight of the object. The area of the circles is proportional to the number of instances. The symbols for the constraints are as follows: u. . .unconstrained, t. . .only translation, r. . .only rotation, x. . .fixed.

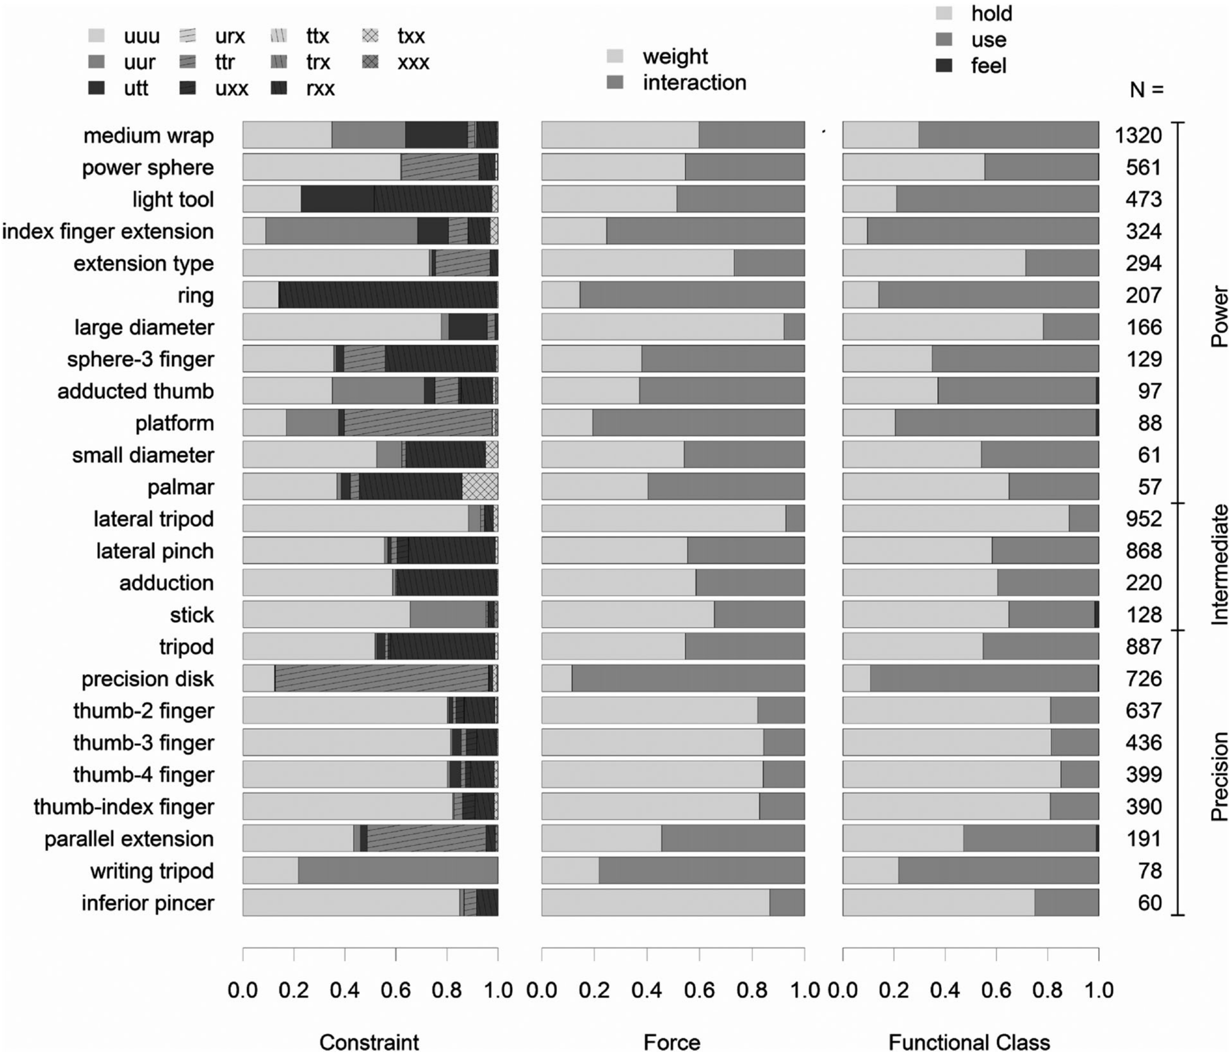

Fig. 5. Task property distributions for each grasp type. The number on the right indicates the number of instances this particular grasp was present in the data set. The left figure shows the task constraints, center shows the force property, and right shows the functional class. Grasps with less than 50 instances are not shown.The order of the legend is equal to the order within each bar.

housekeeping task. In this task, the bottle is lifted, which dominates the grip force requirement, thus “weight” is assigned. Fig. 4 also shows that the functional class is related to the constraints. For the free movement case (“uuu”), most of the instances correspond to “hold. If a constraint is present, most tasks are instead the “use” type.

As is implied in the discussion of task categories described above, there is generally a very strong correlation between the ‘force’ and ‘functional class’ categorization. However, tasks such as carrying a glass of water or pointing a spray bottle deviate from the direct correlation between “force = weight” and “hold”, prompting us to keep both categorizations.

4.3 Correlating Task Properties and Grasp

With regards to the task constraints, in about 54 percent of the instances the object was completely free to move (“uuu”). However, this leaves a full 46 percent of instances where the task is subject to at least one constraint. Concerning the individual grasps (see Fig. 5 left), while unconstrained tasks often dominate, most still exhibit a variety of constraints. One notable exception is the ring grasp, which is mainly used for “rxx”, which is one rotational DoF. This is mainly used for the tasks “turning knob” and “turning handle”. The writing tripod is mainly used in the “urr” case, corresponding to the constraint of one point (the tip of the pen) on a plane.

Concerning the grasp force, Fig. 5 center shows how each single grasped is used. Most of the grasps are dominated by a force related to lifting the object, however, a third of the grasps are used primarily for interaction force tasks. The index finger extension grasp is used mainly for tasks such as mopping and vacuuming; typically constrained tasks. The ring grasp and the palmar grasps are used by the machinist to move controls and covers of machines. Finally, the platform and the precision disk are used mainly for wiping.

As shown in Fig. 5 right, the functional class relationship has similar trends to the force class. The biggest differences occur in the medium wrap, which is used for spray bottles and pliers and the light tool, used primarily for calipers and compressed air nozzles.

4.4 Correlating Task and Object Properties

Up until now this paper has focused on the task-grasp relationship. In the remainder of the paper we combine the task

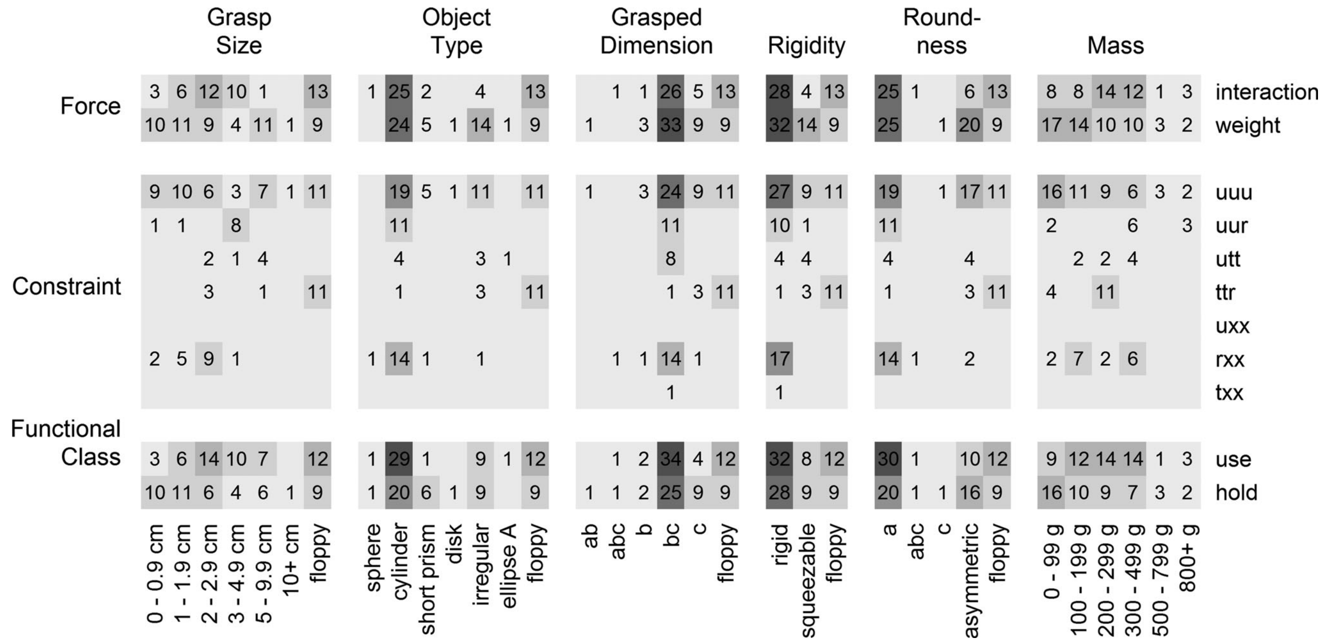

Fig. 6. Relationship between the task and the object properties. Each of the boxes sum up to 100 percent and expresses the relationship between one task attribute and one object attribute. For example taking the top right box, the most common combination is objects with a weight of 0-99 g and a force that is related to the weight of the object, which is the case for 17 percent of all instances. The sum of all squares within the box is 100 percent. Darker background indicates a higher percentage. Rare (smaller than 0.5 percent in each cell) parameters are not plotted. Each box contains at least 98 percent of the instances. The grasp size and the mass were binned as indicated by the labels below. Note that the smallest values from the raters were 1 mm and 1 g.

data with the object data [3], the combined data set contains 7,770 instances. This number is slightly less than the individual object or task data sets, since any instance missing either task or object data was removed.

The relationship between tasks and objects is summarized in Fig. 6. Each of the large boxes represents the relationship between one task and one object attribute, normalized so that the sub-bins sum up to 100 percent. Combinations with a frequency rounded to 0 percent are not plotted to reduce visual clutter. The grasp size and the mass were binned such that the number of instances per bin is well distributed and that the bins make intuitive sense. The number of instances for the grasp size bin is, from small to large: 1,007 (0-0.9 cm), 1,307 (1-1.9 cm), 1,581 (2-2.9 cm), 1,092 (3-4.9 cm), 940 (5-9.9 cm), 140 (10+ cm), and 1,703 floppy objects. For the binning of the mass, the number of instances from small to large is as follows: 1,922 (0-99 g), 1,671 (100-199 g), 1,820 (200-299 g), 1,675 (300-499 g), 246 (500-799 g) and 436 (800+ g).

5 Predictor Quality of Sub-Classifications

In this section, we use machine learning techniques to gain insight into the structure of the data set and to analyze how well the data can predict the grasp type. This process will prove useful in that it will inform not only how well a grasp can be predicted based on task and object data properties as assigned in this study, but also lend insight into how much information is contained within those property classes. As in Section 4.4 above, we use the full data set including task, object, and grasp properties, with a total of 7,770 instances.

A summary of the data set is shown in Table 1. For the object, seven properties are assigned directly and three additional ones are derived based on the assigned parameters. The task data has three attributes and the grasp data has just one attribute—grasp type. All numerical attributes were binned, making them nominal. For the size parameter, the bin size is 1 cm, giving 16 bins, whereas for mass each bin is 100 g wide, resulting in 11 bins. For size C, four bins have no data, thus the final number of levels is 12. Note that for some attributes (A, B, C, Grasp Size) there is a category added for floppy objects, making the actual number of levels one higher. The binning step has two purposes: first, it allows

TABLE 1
Overview of Data Set Attributes

		Attribute	Levels	Dataset 1	2	3	4
Object	Assigned	A	17	x	x	x	
		B	17	x	x	x	
		C	13	x	x	x	
		Grasped Dimension	6	x	x	x	
		Rigidity	4	x	x	x	
		Roundness	5	x	x	x	
		Mass	11	x	x	x	
	Derived	Grasp Size	17	x			
		Shape	5	x			
		Type	12	x			
Task	Assigned	Force	2	x	x		x
		Constraint	11	x	x		x
		Functional Class	3	x	x		x
	Class	Grasp Type	32				

The data set columns on the right indicate the assignment of attributes to the four data sets used in the analysis.

TABLE 2
The Table Reports the Classification Rates in Percent for the Four Data Sets, as Described in Table 1

Dataset	Classification Results
Full (1)	46.6%
Full (2)	46.8%
Object (3)	42.7%
Task (4)	30.8%

The results are based on the decision tree classifier. The nearest neighbor classifier results are within 1.5 percent of the decision tree Classifier.

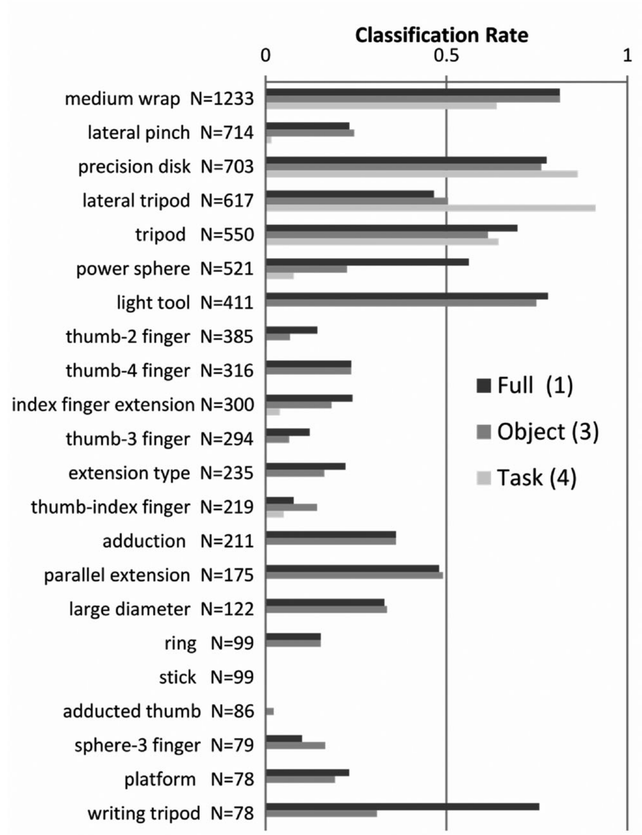

Fig. 7. Classification rates for each grasp type. The results for the Full (1) data set are always within 0.013 of the Full (2) data set, therefore it is not shown in the graph. The grasps are ordered according to their frequency in the data set. Furthermore, grasps with less than 50 instances are not shown. The line at 0.5 corresponds roughly to the classification rate of the full data set (47 percent).

applying classifiers that can only use nominal values; and second, it greatly reduces the chance that the classifier can use the specific numerical values to overfit the data.

In order to investigate the predictive capabilities of our classifications, we consider four combinations of subsets of the classification data. The Full (1) data set contains all attributes from the object (both assigned and derived) and task classifications. With 13 attributes, this is the most comprehensive data set. The Full (2) data set uses only the assigned attributes of object and task, omitting the derived ones, resulting in 10 attributes. Finally, Object (3) and Task (4) use only the object and task attributes, respectively.

In order to examine how well our classifiers predict our data, we apply two machine learning algorithms through the machine learning software Weka [24]. The first classifier we use is a decision tree algorithm, namely the J48 implementation of the C.4.5 algorithm [25], using parameters (-C 0.25, –M 10). The C parameter defines the amount of tree pruning. The minimum number of instances per leaf (parameter -M) was raised to 10 in order to prevent excessive overfitting, since it avoids creating a new leaves for rare parameter combinations. In order to confirm that these results are not an artifact of the particular algorithm used, we also performed a nearest-neighbor classification using the iBK [26] implementation in Weka with parameters (-K 5 –W 0 –I), which define the numbers of neighbors (-K), that the full data set is used for classification (- W 0) and a 1/distance weighting is applied (-I). We found that all results are within 1.5 percent of the decision-tree classification. Furthermore, initial testing with Support Vector Machines using different kernel functions also had similar results. We believe that the close correlation of the results indicates that results described below are based on differences in the data set rather than the classifiers themselves. Due to the fact that we believe a decision tree-based analysis is more logical given the nature of our data, we only report the results from the decision-tree analysis in the following sections.

5.1 Prediction of Grasp Type

In this section, we examine how well each of the four data sets predict grasp type. We use a 10-fold cross-validation [27] to calculate the classification errors. In that procedure the data set is split randomly into 10 equally sized subsamples. Nine subsamples are used to train the classifier and the tenth sample is used to determine the classification accuracy. This step is repeated 10 times, each time alternating which subsample is used for training/testing. The average performance of the classifier is reported. During initial testing, the number of folds used did not appear to significantly influence the classification results, thus we kept the standard 10 fold cross validation.

The classification rates for the four data sets are reported in Table 2. The Full (1) and Full (2) data sets achieve the highest classification rates and the classification rates are basically unchanged by removing the derived parameters to create the Full (2) data. In either case, approximately 47 percent of the grasps are classified correctly. From the individual data sets, the Object (3) properties give higher classification rates (about 43 percent) as compared to the Task (4) data set (about 31 percent). Both classification rates are higher than the baseline classification rate of 16 percent, which would be achieved if all samples were classified as medium wrap, the most common grasp type.

Fig. 7 shows the classification rate for each individual grasp type. For the Full (2) data set only 6 of the 22 most frequent grasp types are above the 0.5 line (corresponding roughly to the overall classification performance of 0.47). The classification rate for the individual Object (3) and Task (4) data sets shows which grasp types can be distinguished well with each type of property. The object set achieves good classification for many grasp types, whereas the task data is generally much less able to distinguish between them. The task classification has a spike on the lateral tripod, which achieves a classification rate of 92 percent. However, in that case that is due to the specifics of the classifier,

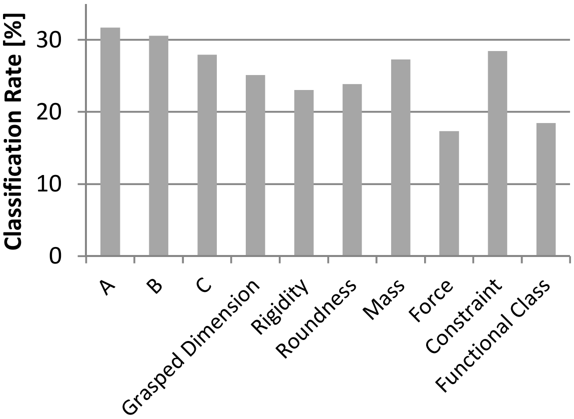

Fig. 8. Classification rate for one attribute based on a decision-tree classifier. All are above the baseline classification rate of 16 percent, which would classify all grasps as medium wrap, the most common grasp type.

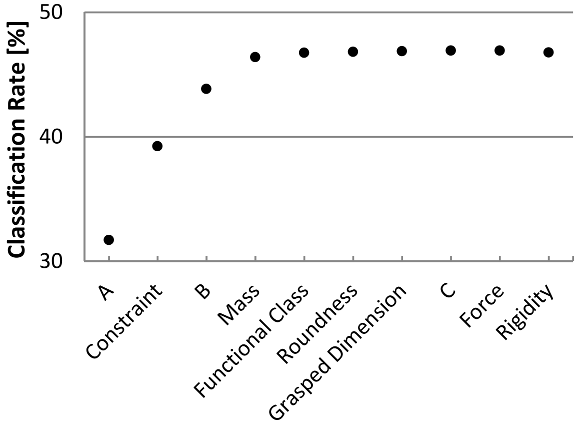

Fig. 9. Classification rates for adding more parameters to the Full(2) data set in the order presented in the figure. A, the most important attribute, is added first. The classification rate plateaus at almost 47 percent.

which classifies 3,456 samples into this category. This means that almost 45 percent of the instances are put into that category, even though they make up for only 8 percent of the data. Consequently, there is a high chance that the real lateral tripod instances are correctly identified.

For some grasp types, the differences in the classification rate for the combined data set (2) and the individual data sets (3 and 4) are relatively high. This indicates that the task and the object classification have complementary information, which can be combined to reliably predict the grasp type. One such case is the writing tripod, which generally has the constraint "uur", requiring that the tip of the pen touches the paper. The writing task is mainly confused with vacuuming, which has the same constraint, but an assignment of medium wrap. For the Object (3) data set, the algorithms cannot distinguish between the adduction grasp and the writing tripod. Thus, only when the object and task data are combined can the result be narrowed down to the final writing tripod grasp. This example shows how combining the object and task data can resolve some ambiguities that would be present if only one source of data was used.

5.2 Importance of Individual Attributes

To estimate how well individual attributes predict the grasp type, we use the classifiers to create an attribute ranking. As the initial results showed that the difference between the Full (1) and Full (2) data sets is very small, we use the Full (2) data for the attribute ranking. This data set does not include the derived properties, yet achieves a similar classification rate.

We apply a greedy search strategy to determine the importance of attributes. We start by calculating the classification rates of the decision tree algorithm for each attribute individually (Fig. 8), with results showing that the size (A, B, and C), constraint, and mass attributes perform best overall. However, these attributes likely contain some redundant information. In order to pick the best combinations of a lower number of attributes, we first select the best attribute and then determine which attribute should be added next in order to maximize the classification rate, repeating this step until we have added all attributes to the data set. Fig. 9 shows the order in which the attributes were added and the gain for each attribute, naturally plateauing at the full classification rate of 47 percent. This strategy does not guarantee an optimal set for a given number of attributes, but will give a set that performs very well. In particular the first attributes are stable, as their initial gain is large. The confidence of the ordering of later attributes diminishes, as the gain of them is very small.

The nearest neighbor classifier agrees on the ordering of the attributes up to mass. For the fifth property it would add the dimension C, whereas the decision tree selects functional class. As the gain for those last attributes is miniscule, this difference is insignificant.

5.3 Prediction of Power, Precision and Intermediate Grasp

To further investigate the prediction capabilities of the data, the classification analysis was performed again with the grasp types replaced by their assignment to the power, intermediate or precision (P/I/P) categories (according to [28]). This results in a data set with 3,308 power grasp instances, 1,674 intermediate grasps and 2,788 precision grasps. Due to the class distribution, the baseline classification rate is 43 percent, which would be achieved if all samples are classified as power grasp.

The results for these data sets are reported in Table 3. As expected, with fewer class levels, the classification rate increased. However, the increase is relatively small, and there is still considerable misclassification. Sixty seven percent classification is achieved with the full data set. The Object (3) data set results in a classification rate of 64 percent and the Task (4) data set achieves 51 percent.

As the reduction in classes is expected to increase the classification rate, a comparison of the results is sought. Therefore, the 32 grasps types in the data set are assigned randomly into three classes and the classification rates are calculated. This step is repeated 100 times and the results are presented in Table 3. The results show that the assignment has a major influence on the final classification results. The best set achieves a classification rate of 77 percent for the Full (1, 2) data set. The number of grasps for each class is 11, 4 and 17, which is similar to the original power (15), intermediate (6) and precision (12) assignment. The mean for the 100 assignments is very similar to the results for the P/I/P data set.

TABLE 3
The Table Reports the Classification Rates for the Data Sets, Where the Grasp Type Was Replaced with Power, Intermediate and Precision [28]

	P/I/P Assignment	Random Assignment		
		Best	Mean	Std
Full (1)	66.6%	77.3%	66.0%	3.5%
Full (2)	66.71%	77.4%	66.0%	3.5%
Object (3)	63.5%	76.2%	63.2%	4.1%
Task (4)	51.1%	72.7%	56.3%	5.1%

The results are contrasted to the classification rates for a random assignment of the grasp types into three categories. The table is based on 100 random assignments, where the results for the best set, the mean and standard deviation for the classification rates is given. The attributes are the same as reported in Table 1 and the results are for the decision tree classifier.

Mass		< 150 g			150 - 350 g			> 350 g		
		Power	Inter-mediate	Precision	Power	Inter-mediate	Precision	Power	Inter-mediate	Precision
Grasp Size [cm]	> 5	45	11	44	77	1	22	89	3	8
	2-5	40	12	48	29	32	39	85	8	6
	< 2	17	36	47	23	43	34	43	31	27

Fig. 10. Distribution of power, intermediate and precision grasp for varying mass and grasp size. The percentages for each mass and grasp size combination sum up to 100 percent. Stronger correlations are highlighted with a darker shading. Floppy objects were excluded, therefore this table is based on 6,067 instances.

While this categorization does not appear to be an effective classifier on average, there may be certain attribute combinations that directly inform the overall type of grasp. To investigate this, we binned the numerical values (size A, B, C) and mass into three levels, to make the resulting rules simple enough to understand and analyze. The bins where chosen to ensure the number of instances within each bin is roughly equal. The same bin thresholds of 2 and 5 centimeters were used for all three size parameters (A, B, and C). Consequently, for the largest size A, most instances are within the "large" category, whereas for C (the smallest object dimension), most instances are in the "small" category. The mass bins are separated by 150 and 350 gram thresholds.

To find association rules within the data, we used the "Associate" function within Weka. In particular, we used the "Apriori" [29] and "Predictive Apriori" [30] algorithms. Most of the more effective association rules predict the power grasp category, largely based on the mass and a size parameter. There are rules that predict precision and intermediate grasps, however, their confidence is much lower and the number of instances is much smaller.

To further investigate the relationship, the frequency of different combinations of size, mass and the grasp class, are shown in Fig. 10. As can be seen, there is a clear trend that heavier and larger objects are more likely to be grasped with a power grasp. For an object with a grasp size larger than 5 cm and a mass over 350 grams, 89 percent of all grasps are power grasps. However, for lightweight and small objects the trend is much less clear, as a small and lightweight object is not necessarily grasped with a precision grasp (47 percent). Power and intermediate grasps are still used to a large degree for small and light objects. Even when the object mass is less than 20 grams, and the grasp size is smaller than 2 cm, the proportion of precision grasps rises only to 61 percent.

6 Discussion

6.1 General Results

The results for task constraints have important implications. Overall, 46 percent of the instances were subject to constraints, which were mainly connected to the functional class "use" and in which the grasp force was usually decoupled from the weight of the object. Simple object transport operations would be adequate only to perform about half of the actions observed. This suggests that for robotic systems, such as assistive robots, designing a robotic system for object transport alone may be inadequate to accomplish the majority of the typical tasks that a human would. It is, however, unclear from the current data set what the typical kinematic differences would be between the object transport and manipulation actions.

In both the object and task classification there are certain classes that have been observed only a few times. In the object case the "fragile" category was present in only 0.2 percent of the instances. This suggests that a high degree of grip force sensitivity in a robotic system may not always be necessary in order to avoid breaking typical objects in human environments, though the machinist objects may be more durable than those typically encountered. Also some grasped dimension combinations (a/b/c, a/b) and roundness classes (c, abc) were found in only a few instances. Regarding tasks, the functional class "feel" was only present in 0.1 percent of the instances. These results suggest that certain object and grasp configurations are quite rare, but it is possible that other professions would encounter some of these rare instances more often.

Some interesting observations can be made using the object and task property distributions in Fig. 6. Concerning the grasp size, it appears that the sizes for force = weight are more evenly distributed than for force = interaction. A similar observation can be made for the functional class—the "use" property is connected to a narrower distribution for grasp size. The main reason for this is likely that all those objects are directly designed to be stably grasped, thus a proper object size was chosen. This fits the results in [31], which claims that the optimal diameter for a handle is 22-32 mm. In general, the distributions for the "use" category can give guidelines for device design, as well as performance specifications for a dexterous robotic hand. If a robotic hand is to be optimized for object transport rather than manipulation, the data from the "hold" category may give more appropriate target specifications or test cases.

Overall, the mass of the object has a relatively uniform distribution up to 500 grams; however, the individual distributions differ depending on the task type. When there is no constraint present ("uuu"), there is a clear trend that the object is lightweight. However, for instances where there is a constraint present, this trend is not observed. For this data set, the participants generally transport lighter objects than they use for a functional purpose. This trend may not hold for all professions—for example a professional mover would likely transport much heavier objects than a machinist or housekeeper.

6.2 Classification Results

Overall, our data set (Full (1) and Full (2)) successfully predicted the grasp type in 47 percent of the instances. The most effective predictor is dimension A (the longest), which outperforms the other dimensions (B and C). This is surprising, since dimension A rarely determines the grasp size directly. One reason for this behavior might be that the longest dimension encodes slightly more information about the object type. As the difference in the classification rate between A and C is 4 percent, slight differences in the data might ultimately lead to this result. The next attributes added are the constraint, size B, and the mass of the objects, each of which add information that increases the classification rate to a larger degree. Using only those attributes, the classification rate is 46 percent—they already capture the majority of the information in the data. These results suggest that for grasp planning systems, major object dimensions, constraint condition, and mass alone could provide much of the information needed to select the overall grasp type.

In 1965, Napier [1] argued that the object shape and task are the two major factors governing how we grasp, however it is still an open question how much shape and task affect the grasp chosen. While this study cannot estimate their full influence, it can help determine a lower bound for the degree of influence. When using only the object data set, about 43 percent of the instances can be classified correctly. For the tasks this number is about 31 percent. Thus, we argue that, as a lower bound, the choice of grasp is influenced at least 43 percent by the object properties and 31 percent by the task properties assigned in this study. Our estimate of task influence may be more of an underestimate, since it is difficult to verbally define exact task constraints, while object properties are more directly related to clear physical quantities. This likely helps explain why the overall task classification rate is lower than for object properties. When the data sets are combined (Full (1) and Full (2)), 53 percent of the instances are incorrectly classified. Much of this remaining error could come from limitations in the rating system. However, in some cases multiple grasps may be equally viable for successful task completion, and humans might randomly select [2] one grasp type.

Our results call into question the conventional classification of grasps into power, intermediate, and precision categories (P/I/P). As they are based on study of grasping [1], [2], one might expect the P/I/P set to perform well. However, our results show that assigning the grasps randomly into three categories results in the same performance. While we might expect that there exists a way to bin grasps into three classes that is easier to classify than the P/I/P categories, it is still surprising that the P/I/P method is not at least easier to classify than a typical random grasp assignment. It is possible that the P/I/P classification would perform better with a different selection of objects, such as by including more heavy objects.

Currently, the only reliable prediction about the P/I/P assignment was for power grasps. The results in Section 5.3 show that for heavy (>350 g) and large (>5 cm) objects, 89 percent of all grasps were power grasps. However, the opposite is not true. One reason for this asymmetry might be the fact that a large and/or heavy object must be grasped by a grasp that can handle such an object. The opposite is not true, a small object can still be grasped by a power or intermediate grasp. This result is quite important, as it provides evidence against a commonly held notion that power grasps are used mainly for heavier objects or in higher force tasks. Our data suggests instead that small objects may also benefit from the apparent added security of a power grasp. For robotic applications, this suggests that designing hands to allow power grasps of smaller objects may be beneficial, and also that grasp planners should not necessarily pick precision grasps for small objects.

7 Conclusions and Future Work

We set out to augment the data from our video analysis study [4] with information about the objects being manipulated and the tasks being performed. The data shows grasping behavior of two machinists and two housekeepers being recorded during their professional work; analyzing other professions could add further generalizability to the results. The data highlights the complexity of human manipulation. With a set of seven object and three task properties we could predict the grasp type correctly in about 47 percent of the instances. The upper limit of the possible classification rate in this study is the inter-rater agreement of the grasp assignments, which was 62 percent for our data set [4]. Although greater rater agreement could improve this statistic somewhat, we believe that human manipulation is sufficiently complex and diverse such that there will always be a substantial level of uncertainty.

Certain major results seem particularly useful or important. Many (46 percent) of instances observed are subject to constraints, which emphasizes that many real-world human tasks are not as simple as object transport. It is thus worthwhile to design robotic systems for human environments to allow more versatile behavior. The classification results suggest that major object dimensions, constraint condition, and mass are particularly important to selecting grasp type, and these criteria could thus be used as effective heuristics for grasp planning systems. Finally, the results call into question the conventional notions of power and precision grasps, and suggest that while power grasps are well suited to heavy objects or high force tasks, they may still be useful even for very small or lightweight objects. This suggests robotic systems could still take advantage of power grasps even when interacting with small objects.

Our current approach in this and related papers [3], [20] provides useful information about the objects that we interact with, the tasks we perform, and the grasps we use for them. These results can be applied to define performance specifications and test conditions for robotic hands, to provide basic heuristics useful in developing grasp planners, to target rehabilitation efforts toward essential hand functionality, and finally, to aid in the design of haptic interfaces or other devices that should interact with the hand in a natural manner.

ACKNOWLEDGMENTS

The authors would like to thank Charlotte Guertler, Joshua Zheng, and Sara De La Rosa for their work in coding the video data set used in this paper. This work was supported in part by a grant from the US National Science Foundation (NSF), Grant NSF IIS- 0952856.

REFERENCES

[1] J. R. Napier, "The prehensile movements of the human hand," *J. Bone Joint Surg. Brit. Vol.*, vol. 38-B, no. 4, pp. 902–913, Nov. 1956.
[2] N. Kamakura, M. Matsuo, H. Ishii, F. Mitsuboshi, and Y. Miura, "Patterns of static prehension in normal hands," *Amer. J. Occupational Therapy: Official Publication Amer. Occupational Therapy Assoc.*, vol. 34, no. 7, pp. 437–445, Jul. 1980.
[3] T. Feix, I. M. Bullock, and A. M. Dollar, "Analysis of human grasping behavior: Object characteristics and grasp type," *IEEE Trans. Haptics*, 2014.
[4] I. M. Bullock, J. Z. Zheng, S. D. La Rosa, C. Guertler, and A. M. Dollar, "Grasp frequency and usage in daily household and machine shop tasks," *IEEE Trans. Haptics*, vol. 6, no. 3, pp. 296–308, Jul.-Sep. 2013.
[5] I. M. Bullock, T. Feix, and A. M. Dollar. (2014). The Yale human grasping data set: Grasp, object and task data in household and machine shop environments [Online]. Available: http://www.eng.yale.edu/grablab/humangrasping/
[6] S. Katz, R. Moskowitz, J. Ba, and M. Jaffe, "Studies of illness in the aged-the index of ADL: A standardized measure of biological and psychosocial function,"*J. Amer. Med. Assoc.*, vol. 185, no. 12, pp. 914–919, 1963.
[7] I. Hartigan, "A comparative review of the Katz ADL and the Barthel Index in assessing the activities of daily living of older people," *Int. J. Older People Nursing*, vol. 2, no. 3, pp. 204–12, Sep. 2007.
[8] M. Philipose, K. Fishkin, M. Perkowitz, D. Patterson, D. Fox, H. Kautz, and D. Hähnel, "Inferring activities from interactions with objects," *IEEE Pervasive Comput.*, vol. 3, no. 4, pp. 50–57, Oct. 2004.
[9] International classification of functioning, disability and health: ICF, World Health Organization, Genève, Switzerland, 2001.
[10] Y. Matsumoto, Y. Nishida, Y. Motomura, and Y. Okawa, "A concept of needs-oriented design and evaluation of assistive robots based on ICF," *Proc. IEEE Int. Conf. Rehabil. Robot.*, 2011, pp. 1–6.
[11] M. Ehling, D. Schäfer, I. Kahle, I. Meyer, B. Wende, C. Wingerter, and R. Wilhelm, "Wo bleibt die Zeit? Die Zeitverwendung der Bevölkerung in Deutschland 2001/02," Statistisches Bundesamt, Bundesministerium für Familie, Senioren, Frauen und Jugend, 2003.
[12] "How Australians use their time," Australian Bureau of Statistics, Canberra, A.C.T., Australia, Catalogue Number: 4153.0, 2006.
[13] R. L. Klatzky, B. McCloskey, S. Doherty, J. Pellegrino, and T. Smith, "Knowledge about hand shaping and knowledge about objects," *J. Motor Behavior*, vol. 19, no. 2, pp. 187–213, 1987.
[14] C. Hesse and H. Deubel, "Advance planning in sequential pick—and—Place tasks," *J. Neurophysiol.*, vol. 104, no. 1, pp. 508–516, Jul. 2010.
[15] S. Johnson Frey, M. McCarty, and R. Keen, "Reaching beyond spatial perception: Effects of intended future actions on visually guided prehension," *Vis. Cognition*, vol. 11, nos. 2/3, pp. 371–399, Mar. 2004.
[16] R. Cohen and D. Rosenbaum, "Where grasps are made reveals how grasps are planned: generation and recall of motor plans," *Exp. Brain Res.*, vol. 157, no. 4, pp. 486–495, Aug. 2004.
[17] C. Ansuini, L. Giosa, L. Turella, G. Altoè, and U. Castiello, "An object for an action, the same object for other actions: effects on hand shaping," *Exp. Brain Res.*, vol. 185, no. 1, pp. 111–119, Feb. 2008.
[18] C. Ansuini, M. Santello, S. Massaccesi, and U. Castiello, "Effects of end-goal on hand shaping," *J. Neurophysiol.*, vol. 95, no. 4, pp. 2456–2465, Apr. 2006.
[19] P. Haggard, "Planning of action sequences," *Acta Psychol.*, vol. 99, no. 2, pp. 201–215, Jul. 1998.
[20] J. D. Morrow and P. K. Khosla, "Manipulation task primitives for composing robot skills," in *Proc. IEEE Int. Conf. Robot. Autom.* , 1997, vol. 4, pp. 3354–3359.
[21] G. Morris and L. Haynes, "Robotic assembly by constraints," in *Proc. IEEE Int. Conf. Robot. Autom.*, 1987, pp. 1507–1515.
[22] J. Z. Zheng, S. D. La Rosa, A. M. Dollar, and S. De La Rosa, "An investigation of grasp type and frequency in daily household and machine shop tasks," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2011, pp. 4169–4175.
[23] A. J. Conger, "Kappa reliabilities for continuous behaviors and events," *Edu. Psychol. Meas.*, vol. 45, no. 4, pp. 861–868, Dec. 1985.
[24] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. Witten, "The WEKA data mining software: An update," *ACM SIGKDD Explorations*, vol. 11, no. 1, pp. 10–18, 2009.
[25] S. Salzberg, "C4.5: Programs for Machine Learning by J. Ross Quinlan. Morgan Kaufmann Publishers, Inc., 1993," *Mach. Learn.*, vol. 16, no. 3, pp. 235–240, 1994.
[26] D. Aha, D. Kibler, and M. Albert, "Instance-based learning algorithms," *Mach. Learn.*, vol. 6, no. 1, pp. 37–66, 1991.
[27] R. Kohavi, "A study of cross-validation and bootstrap for accuracy estimation and model selection," in *Proc. 14th Int. Joint Conf. Artif. Intell.*, 1995, pp. 1137–1143.
[28] T. Feix, R. Pawlik, H.-B. Schmiedmayer, J. Romero, and D. Kragic, "A Comprehensive Grasp Taxonomy," in *Robot., Sci. Syst.: Workshop Understanding Human Hand Advancing Robot. Manipulation*, 2009.
[29] R. Agrawal and R. Srikant, "Fast algorithms for mining association rules," in *Proc. 20th Int. Conf. Very Large Data Bases*, 1994, vol. 1215, pp. 487–499.
[30] T. Scheffer, "Finding association rules that trade support optimally against confidence," in *Principles of Data Mining and Knowledge Discovery*, vol. 2168, L. Raedt and A. Siebes, Eds. Berlin, Heidelberg, Germany: Springer, 2001, pp. 424–435.
[31] A. Tilley and H. Associates, *The Measure of Man and Woman: Human Factors in Design*. Hoboken, NJ, USA: Wiley, 2001.

Thomas Feix received the MSc degree in sports equipment technology from the University of Applied Sciences Technikum Wien, Vienna, Austria, and the PhD degree from the Vienna University of Technology in 2011. He is a postdoctoral associate with the GRAB Lab, Department of Mechanical Engineering, Yale University, New Haven, Connecticut. His research is focused on human grasping and manipulation, and its application to robotics and prosthetics. He is a member of the IEEE.

Ian M. Bullock (S'11) received the BS degree in engineering from Harvey Mudd College, Claremont, California, in 2010. His work experience has ranged from programming educational neuroscience modules to designing circuits for a particle accelerator. He is working toward the PhD degree at Yale University since Fall 2010. His research interests include human dexterity and robotic manipulation, with particular emphasis on using an understanding of the capabilities of the human hand to improve the design of robotic and prosthetic manipulators. He is a student member of the IEEE.

Aaron M. Dollar received the BS degree in mechanical engineering from the University of Massachusetts at Amherst, and the SM and PhD degrees in engineering sciences from Harvard University. He conducted two years of postdoctoral research at the MIT Media Lab. He is the John J. Lee associate professor of mechanical engineering and materials science at Yale University. His research topics include human and robotic grasping and dexterous manipulation, mechanisms and machine design, and assistive and rehabilitation devices including upper-limb prosthetics and lower-limb orthoses. He received the 2013 DARPA Young Faculty Award, 2011 AFOSR Young Investigator Award, 2010 Technology Review TR35 Young Innovator Award, and the 2010 US National Science foundation (NSF) CAREER Award. He is a member of the IEEE.